On constructive recognition
of a black box $\mathrm{PSL}(d, q)$*

Peter A. Brooksbank and William M. Kantor

Abstract. Assuming the availability of an oracle for handling black box groups isomorphic to $\mathrm{SL}(2, q)$, we present an algorithm that constructively recognises a group G known to be isomorphic to $\mathrm{PSL}(d, q)$ for known $d \geq 3$ and q. We also outline an analogous result for symplectic groups.

1. Introduction

The presently known algorithms for constructive recognition of black box classical simple groups have asymptotic running times that are not polynomial **[Br, KS]**: the complexity contains factors of q (the field size), whereas the input size involves only $\log q$. In this paper, we will take a step toward the polynomial time paradigm for the $\mathrm{PSL}(d, q)$ case, removing all factors of q in the timing at the expense of making calls to an oracle that allows us to work inside subgroups isomorphic to $\mathrm{SL}(2, q)$; thus, $\mathrm{SL}(2, q)$ is the polynomial-time bottleneck.

Consequently, we pay close attention to the cost χ of each call to our hypothesised $\mathrm{SL}(2, q)$-oracle, in addition to the more standard timing parameters used in **[KS]** (namely, the input length, which is greater than $(d^2 \log q)/2$, together with the time μ required to perform group operations in G, and the time ξ required per element for the construction of independent, (nearly) uniformly distributed random elements of G). The motivation for this oracle hypothesis is provided by a recent advance by Leedham-Green and Conder reported at this conference **[CLG]**: an algorithm that deals with the groups $\mathrm{SL}(2, q)$ as 2×2 matrix groups assuming the availability of an oracle for the discrete log problem. This advance prompted Leedham-Green to ask us the question answered by the following result (where SLP stands for "straight-line program"; cf. **[KS]**, 2.2.5):

Theorem 1.1. *Let $G = \langle \mathcal{S} \rangle$ be a black box group, believed to be isomorphic to a nontrivial homomorphic image of $\mathrm{SL}(d, q)$ for known $q \geq 17$ and $d \geq 3$, and equipped with an oracle for handling black box groups isomorphic to $\mathrm{SL}(2, q)$. Then the following hold.*

(1) *There is a Las Vegas algorithm which, with probability $\geq \frac{1}{2}$, in time polynomial in ξ, χ, μ, d and $\log q$, namely in*

$$O(d^2 \log(d \log q) \log q \{\xi + \chi d \log q + \mu d^2 \log^2 q\}$$
$$+ |\mathcal{S}| \log(d|\mathcal{S}|)\{\xi + \chi d^2 \log q + \mu d^2 \log^2 q + d^5 \log q\})$$

time, verifies that G is a homomorphic image of $\mathrm{SL}(d, q)$ and constructs a new generating set $\mathcal{S}^*$ for G, a generating set X for $\mathrm{SL}(d, q)$, and a bijection $\Psi : X \to \mathcal{S}^*$ that extends to an epimorphism $\Psi : \mathrm{SL}(d, q) \to G$.

(2) *X, $\mathcal{S}^*$ and Ψ have the following properties.*

    (a) *In $O(d\chi + \mu d^2 \log q)$ time, when given $A \in \mathrm{SL}(d, q)$ one can construct an SLP of length $O(d^2 \log q)$ from X to A, and then mimic it in G to obtain $\Psi(A)$.*

    (b) *In Las Vegas $O(\log d\{\xi + \chi(d + \log q) + \mu d \log^2 q\} + d^5 \log q)$ time, with probability $\geq 1 - \frac{1}{8d^2}$ an SLP of length $O(d^2 \log q)$ can be found from $\mathcal{S}^*$ to any given $g \in G$, and then mimicked in $\mathrm{SL}(d, q)$ to obtain $\Psi^{-1}(g)$ modulo scalars.*

Note that failure in (1) could be due to bad luck with random selections, or to the fact that G is not isomorphic to the stated type of quotient group; but if there is an output then it is guaranteed to be correct (since the algorithm is Las Vegas). The lower bound on q is designed to simplify our arguments; the algorithm of **[KS]** applies for the remaining small q. This paper is intended, primarily, as a commentary on the $\mathrm{PSL}(d, q)$ algorithm presented in **[KS]**. We frequently refer to results and routines contained therein with little further comment, and will generally adhere to notational conventions used there.

The case $d = 3$ is the hardest one we consider (cf. Section 3). The general case, in Section 2, needs only a few non-obvious modifications of **[KS]**. We note that we use our $\mathrm{SL}(2, q)$-oracle for various subgroups isomorphic to $\mathrm{SL}(2, q)$, many of which involve somewhat random generators; it would have been far better to have used the oracle only a few times. However, it is hard to imagine not needing at least $|\mathcal{S}|$ oracle calls in order to verify that $\langle \mathcal{S} \rangle = \langle \mathcal{S}^* \rangle$.

We have only dealt with homomorphic images of $\mathrm{SL}(d, q)$, but homomorphic images of other subgroups of $\mathrm{Aut}\,\mathrm{SL}(d, q)$ containing $\mathrm{SL}(d, q)$ can be reduced to this case using a Monte Carlo algorithm for finding the derived group **[BCFLS]**.

In Section 4 we outline an analogue of this theorem for symplectic groups and comment on the other classical groups. We have continued to follow the recursive approach in **[KS]**. There is some hope that this can be avoided. In particular, we

suspect that the methods used in **[Br]** for the group PSL(d,q) also can be modified in the same way that **[KS]** has been in this paper.

The SL$(2,q)$-oracle.   We assume that our hypothesised oracle is able to perform various tasks in a black box group $L = \langle \mathcal{S} \rangle$ (believed to be) isomorphic to SL$(2,q)$ for known q. We list these tasks together with the time required to perform them.

(i) The oracle provides us with sets $\mathcal{S}^* \subset L$ and $X = \left\{ \begin{pmatrix} 1 & \rho^k \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ \rho^k & 1 \end{pmatrix} \mid 0 \leq k < e \right\} \subset$ SL$(2,q)$ for a known generator ρ of GF$(q)^*$, and a bijection $\Psi : X \to \mathcal{S}^*$, behaving as in Theorem 1.1. We denote by χ the time required for Theorem 1.1(1) as well as for each application of Theorem 1.1(2).

(ii) The oracle is also equipped with routines for discrete logarithms in GF$(q)^*$. Let T denote the transvection group of L generated by the image under Ψ of $\left\{ \begin{pmatrix} 1 & 0 \\ \rho^k & 1 \end{pmatrix} \mid 0 \leq k < e \right\}$ for ρ in (i), and suppose we are given $t \in T$. We can use the oracle to find $\Psi^{-1}(t) = \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}$ for some $\lambda \in$ GF(q). In addition, the oracle can compute the integer n such that $0 \leq n < q - 1$ and $\lambda = \rho^n$. We assume that the cost per call to the discrete log oracle is χ.

We will find task (ii) particularly useful in the following algorithmic setting. If we have elements $g, h \in N_G(T)$ (but not necessarily in L) both inducing scalars of the same order on T, then h can be replaced by a specific power of itself to ensure that g and h induce the same scalar on T. The discrete log oracle tells us which power.

We view this oracle as a black box. However, in practice it will presumably be replaced by a Las Vegas algorithm for Theorem 1.1(1): the algorithms in **[KS]** and **[CLG]** for this is are, indeed, Las Vegas. On the other hand, we assume that our discrete log oracle is deterministic in practice.

2. The general case: $d \geq 4$

Somewhat as in **[KS]**, the methods we employ for dimension 3 differ widely from the larger dimensions. In this section, we assume that $d \geq 4$, and whenever possible we follow closely the algorithm in **[KS]**. In particular, the algorithm we present here is recursive. We wish to ensure that the probability of the routine failing up to the point where the recursive call is made in Section 2.6 is no more than $\frac{1}{d^2}$. In any non-deterministic step in the main routine, this is achieved by making an extra factor of (at least) $\lceil \log d \rceil$ random choices.

2.1. Background. We recall some background material discussed at greater length in **[KS]**.

(1) Groups of transvections. We assume a familiarity with the basic properties of transvections and transvection subgroups of $\mathrm{SL}(d, q)$. For a hyperplane H of V and a point $\alpha \in H$, we denote by $T_{\alpha, H}$ the group of transvections with centre α and axis H. As in **[KS]**, we denote by $Q(\alpha)$ the elementary abelian group of order q^{d-1} consisting of all transvections with centre α, and by $Q = Q(H)$ the group of all transvections with axis H. In this way, we may view points (resp. hyperplanes) as existing within the group G as subgroups $Q(\alpha)$ (resp. $Q(H)$).

(2) Primitive prime divisors. By a fundamental theorem of Zsigmondy **[Zs]**, if p is prime and $m \geq 2$ then there is a prime dividing $p^m - 1$ but not $p^i - 1$ for $1 \leq i < m$, except when either $p = 2, m = 6$, or $m = 2$ and p is a Mersenne prime. Such a prime is called a *primitive prime divisor* of $p^m - 1$. We will call an integer j a ppd$^{\#}(p; m)$-*number* if $j | p^m - 1$ and either $p = 2, m = 6$ and $21 | j$; $m = 2$, p is a Mersenne prime, and $8 | j$; $m = 1$ and $j > 4$; or j is divisible by a primitive prime divisor of $p^m - 1$ (this is slight departure from **[KS]**, taking advantage of the fact that here we will always have $p^m \geq 7$).

   We test whether the order of an element of G is a ppd$^{\#}(p; m)$-number using a deterministic $O(\mu m^2 \log p)$ time algorithm in **[KS]**, 2.6.

(3) Field computations. These are discussed in **[KS]**, 2.3. We will have at our disposal the hypothesised $\mathrm{SL}(2, q)$ and discrete log oracles in place of the Zech table used there.

2.2. Time involving factors of q in [KS]. Since our goal here is only to remove occurrences of factors of q in the timing in **[KS]**, we begin by pinpointing those occurrences, which were of five types when $d \geq 4$:

   (i) finding certain elements of G whose probability of occurrence is bounded below only by $\frac{1}{q^M}$ for some integer M;

   (ii) testing for isomorphism to $\mathrm{SL}(3, q)$ and finding such an isomorphism;

   (iii) performing a test using all conjugates of the form $A^b, b \in B$, for a transvection group B and an element or subgroup A;

   (iv) listing a full transvection group of order q, primarily in order to perform linear algebra within that group; and

   (v) a discrete log computation in $\mathrm{GF}(q)^*$ (using a Zech table for $\mathrm{GF}(q)$).

Type (i). This was used to obtain transvections when $d \geq 4$ (**[KS]**, 3.2.1), and here is replaced by Sections 2.3 for $d \geq 5$ and 2.4 when $d = 4$.

   In addition to finding transvections, by using the probability estimate in (i), **[KS]**, 3.2.1, obtained an element of suitable ppd order that is not available here when d is even. Hence we use substitute procedures to find certain subgroups Q, $Q(\alpha)$, H, L and $\mathcal{S}^*$ when $d \geq 5$ (cf. **[KS]**, 3.3.1). However, in the present situation we only find these subgroups by a Monte Carlo algorithm (in Section 2.5), verifying correctness later in Section 2.8 and hence only then converting this to a Las Vegas algorithm.

Type (ii). In Section 3 we give such a test.

Type (iii). This occurred only in **[KS]**, Lemma 3.11 (effective transitivity of the groups Q and $Q(\alpha)$), which here is replaced by Lemma 2.1. This transitivity was then used several times. First it was used in **[KS]**, Corollary 3.12 to find a subgroup $L \cong \mathrm{SL}(d-1, q)$; this is replaced in Section 2.5. Then it was used in **[KS]**, 3.5.1, to label the points of the target vector space, which in turn was needed in **[KS]**, 3.5.2, to construct the desired homomorphism $G \to \mathrm{PSL}(d, q)$; see Section 2.8. Also it was used in **[KS]**, Proposition 3.18, for the analogue of our Theorem 1.1(2b); again see Section 2.8.

Type (iv). This occurred in **[KS]**, 3.4.3(1),(4),(5): finding and computing with GF(p)- and GF(q)-bases of Q and $Q(\alpha)$; expressing any element of Q or $Q(\alpha)$ as linear combinations or SLPs using these bases; and finding matrices of the linear transformations induced on Q or $Q(\alpha)$. See Section 2.7.

Type (v). This occurred in **[KS]**, Propositions 3.17 and 3.18, where a discrete log calculation ensured that a certain $(d-1) \times (d-1)$ matrix had determinant 1. Of course our discrete log oracle is used here.

[KS], 7.3.2. Verifying a presentation is the only other place where a factor q appears in a timing. This verification uses all of the above procedures, and these already have had q removed from their timings.

2.3. Constructing transvection groups and J when $d \geq 5$.

Calls to $d = 3$ or 4. We will make frequent use of the cases $d = 3$, 4 of Theorem 1.1 for subgroups generated by very special subgroups of G, although we will not use the additional time required to verify a presentation. With small probability, the subgroups generated by our chosen elements might not be SL$(3, q)$ (or SL$(4, q)$) and the failure goes undetected. However, in this case, the algorithm will eventually fail, and we will not succeed in our goal of finding and verifying a presentation for G. We will see, in Sections 2.8 and 3, that a call to Theorem 1.1(1), without the additional presentation verification, costs $O(\xi + \chi e + \mu \log^2 q)$ when $d = 3$ and $O(e \log e(\xi + \chi e + \mu \log^2 q))$ when $d = 4$.

The elements a_1, a_2. We first present an $O(d \log d(\xi + \mu e d^2 \log q))$-time Las Vegas algorithm which, with probability $\geq 1 - \frac{1}{8d^2}$, constructs an element $a \in G$ which, in its action on the sought-after vector space, has 2-dimensional support and is the identity on a $d - 2$-space (for even d we construct two such elements). We also construct $\sigma \in G$ centralising a and having $d - 3$- or $d - 2$-dimensional support.

Let $r = d - 2$ for d odd and $d - 3$ for d even. Choose up to $\lceil 32(d-2) \log(2d) \rceil$ elements τ of G, and for each test whether $|\tau|$ is divisible by both a ppd$^{\#}(p; er)$- and ppd$^{\#}(p; 2e)$-number. If so let $a = \tau^{q^r - 1}$ and $\sigma = \tau^{q^2 - 1}$. Then $|a|$ is a ppd$^{\#}(p; 2e)$-divisor of $q + 1$.

Let $\tau_1 = \tau_2 = \tau$ if d is odd. When d is even, repeat the above a second time in order to obtain two pairs (a_1, σ_1) and (a_2, σ_2), where $a = a_1$.

The proofs of correctness, timing and probability of success of this procedure are similar to those in **[KS]**, 4.2.1 and 5.2.1. Note that no oracle calls are needed here.

Constructing $K \cong$ SL(4, q). We next present an $O((\log d) \cdot (e \log e)(\xi + \chi e + \mu \log^2 q))$-time Las Vegas algorithm which, with probability $\geq 1 - \frac{1}{8d^2}$, constructs a naturally embedded subgroup $K \cong \mathrm{SL}(4, q)$ together with a generating set $\mathcal{S}_K^*$ of K, a field $F = \mathrm{GF}(q)$, a vector space V_K and an isomorphism $\mathrm{SL}(V_K) = \mathrm{SL}(4, q) \to K$ behaving as in Theorem 1.1.

Choose up to $\lceil 16 \log(2d) \rceil$ elements $g \in G$, and for each use the case $d = 4$ of Theorem 1.1 to test whether $K = \langle a_1, a_2^g \rangle \cong \mathrm{SL}(4, q)$, in which case also obtain an isomorphism $\Psi_K : \mathrm{SL}(4, q) \to K$.

Note that, in view of the subgroup structure of $\mathrm{SL}(4, q)$, two elements having 2-dimensional supports and whose orders are $\mathrm{ppd}^\#(p; 2e)$-divisors of $q + 1$ generate $\mathrm{SL}(4, q)$ with probability at least $\frac{1}{2}$ (this result is the concatenation of many papers, summarized in **[KL]**, Theorem 5.1). The stated timing arises from the $\Theta(\log d)$ calls to the case $d = 4$ of the theorem.

At this point we have a field $F = \mathrm{GF}(q)$; fix a generator ρ of F^* throughout the remainder of the proof (observe that this generator is available from just one call to the $\mathrm{SL}(2, q)$-oracle).

Constructing $J \cong$ SL(3, q). We prefer, for two reasons, to work with a subgroup $\mathrm{SL}(3, q)$ rather than the $\mathrm{SL}(4, q)$ that we have constructed. Firstly, using an $\mathrm{SL}(3, q)$ subgroup matches up better with **[KS]**, making it easier for the reader to translate. Also, we will need to construct an $\mathrm{SL}(3, q)$ subgroup when we deal with $\mathrm{SL}(4, q)$ in Section 2.4.

Therefore we now construct a subgroup $J \cong \mathrm{SL}(3, q)$ of K, succeeding with probability $\geq 1 - \frac{1}{8d^2}$.

Use Theorem 1.1(2b) $O(\log d)$ times for Ψ_K to find matrices $\Psi_K^{-1}(a_1)$ and $\Psi_K^{-1}(a_2^g)$ in $\mathrm{SL}(V_K)$. Use linear algebra to compute the 2-spaces $W_1 = [V_K, \Psi_K^{-1}(a_1)]$ and $W_2 = [V_K, \Psi_K^{-1}(a_2^g)]$, and find $A \in \mathrm{SL}(V_K)$ such that $W_1 \cap W_2 A$ is 1-dimensional. Use Theorem 1.1(2a) to find $\Psi_K(A)$. Replace $a_2 \leftarrow a_2^{g \Psi_K(A)}$, $\sigma_2 \leftarrow \sigma_2^{g \Psi_K(A)}$ and set $J = \langle a_1, a_2 \rangle$.

Next, use Section 3 $O(\log d)$ times to find a generating set $\mathcal{S}_J^*$ of J, and an isomorphism $\Psi_J : \mathrm{SL}(V_J) \to J$ behaving as in Theorem 1.1. Use Theorem 1.1(2b) $O(\log d)$ times for Ψ_J to find the matrices $\Psi_J^{-1}(a_1)$ and $\Psi_J^{-1}(a_2)$ in $\mathrm{SL}(V_J)$.

The purpose of each of the $O(\log d)$ repetitions above is to ensure that we correctly find all of the matrices $\Psi_K^{-1}(a_1)$, $\Psi_K^{-1}(a_2^g)$, $\Psi_J^{-1}(a_1)$, $\Psi_J^{-1}(a_2)$ and the isomorphism Ψ_J with probability $\geq 1 - \frac{1}{8d^2}$. All of the remaining uses for Ψ_J will be to compute images of matrices in the group J, using the deterministic Theorem 1.1(2a), so there will be no more randomness needed in this section.

Some elements and subgroups of J. Use linear algebra in V_J to compute the point α of intersection of the 2-spaces $V_{2,i} = [V_J, \Psi_J^{-1}(a_i)]$, $i = 1, 2$. Find $B, C \in \mathrm{SL}(3, q)$

such that $T^B_{\alpha, V_{2,2}} = T_{\alpha, V_{2,1}}$ and $C = \mathrm{diag}(1, \rho, \rho^{-1})$, using a basis of V_J starting with a vector in α.

Use Ψ_J to find the following: $f = \Psi_J(B)$, $T = \Psi_J(T_{\alpha, V_{2,1}})$ (centralised by both σ_1 and σ_2^f); $j = \Psi_J(C)$ (inducing on T an automorphism of order $q - 1$); the two subgroups X_1, X_2 of J having order q^2, consisting of transvections and containing T; any $j(\gamma) \in J$ such that $[[X_1, X_2^{j(\gamma)}], X_1] \neq 1$ (i.e., X_1 and $X_2^{j(\gamma)}$ are not "incident" within J); and the subgroup $D \cong \mathrm{SL}(2, q)$ of J normalising both X_1 and $X_2^{j(\gamma)}$. For each of the above elements of J we have an $O(\log q)$-length SLP from $\mathcal{S}_J^*$, found in the time $O(\chi + \mu \log q)$ to use Theorem 1.1(2a) for J.

The total time in Section 2.3 is $O(d \log d(\xi + \mu e d^2 \log q) + (\log d)(e \log e)(\xi + \chi e + \mu \log^2 q))$, and it succeeds with probability at least $1 - \frac{1}{4d^2}$.

2.4. Constructing transvection groups and J when $d = 4$.

Next we also construct a naturally embedded subgroup $J \cong \mathrm{SL}(3, q)$ when $d = 4$.

As in Section 2.3, find an element τ whose order order is divisible by both a ppd$^\#(p; e)$- and ppd$^\#(p; 3e)$-number, and set $a = \tau^{3(q^2+q+1)}$. Consider conjugates a^{g_1}, a^{g_2}, let $A = \langle a, a^{g_1}, a^{g_2} \rangle$, find the derived subgroup $J = A'$ using **[BCFLS]**, and use Section 3 to test whether $J \cong \mathrm{SL}(3, q)$. If this isomorphism holds, we obtain a field, an isomorphism $\Psi_J : \mathrm{SL}(3, q) = \mathrm{SL}(V_J) \to J$ and a generating set $\mathcal{S}_J^*$ for J.

By an argument similar to that in **[KS]**, Lemma 3.7, we have $A \cong \mathrm{SL}(3, q) \circ \langle a \rangle$ with probability $\geq \frac{1}{2^5}$, while Section 3 succeeds with probability $\geq \frac{1}{2}$. Hence, we obtain J and then Ψ_J with probability $\geq \frac{1}{2^6}$. Finding A' uses the $O(\mu \log^2 q)$-time black box Monte Carlo algorithm in **[BCFLS]**.

Note that there are no elements σ, σ_i here; and f will not be needed. Use the isomorphism Ψ_J and linear algebra to find the following: the 2-dimensional eigenspace H and the 1-dimensional eigenspace γ of $\Psi_J^{-1}(a)$ in V_J; the transvection group $T = \Psi_J(T_{\alpha, H})$ for some point $\alpha \in H$; $B' \in \mathrm{SL}(V_J)$ sending α to γ and $j(\gamma) = \Psi_J(B')$; a diagonal matrix C inducing ρ on $T_{\alpha, H}$ and $j = \Psi_J(C)$; and X_1, X_2, D as in Section 2.3. This takes $O(\xi + \chi e + \mu \log^2 q)$ time.

2.5. The subgroups Q, $Q(\alpha)$, $Q(\gamma)$, H and L.

If $d \geq 5$, let

$$Q = \langle X_1^{\sigma_1^i}, X_1^{(\sigma_2^f)^i} \mid 0 \leq i \leq d - 2 \rangle, \quad Q(\alpha) = \langle X_2^{\sigma_1^i}, X_2^{(\sigma_2^f)^i} \mid 0 \leq i \leq d - 2 \rangle$$

and $Q(\gamma) = Q(\alpha)^{j(\gamma)}$. We claim that, with probability $> 1 - \frac{4}{q^{d-4}}$, Q and $Q(\alpha)$ are the subgroups of G, of order q^{d-1}, consisting of all transvections having the same centre (or axis) as T (cf. **[KS]**, 3.1.3). We will only prove this for the group Q and the case of even $d \geq 6$. Up to duality in the target vector space V, the stated conjugates of X_1 all have the same axis W. With probability $> 1 - \frac{2}{q^{d-4}}$ the codimension 1 or 2 subspaces of W preserved by $\sigma_1, \sigma_2^f \in C_G(T)$ generate W, in which case

$\langle X_1^{\langle \sigma_1 \rangle}, X_1^{\langle \sigma_2^f \rangle} \rangle$ is the desired group of order q^{d-1}. Finally, the stated conjugates of X_1 generate these normal closures by **[KS]**, Lemma 2.7.

We emphasise that *we do not* yet *know for certain that the output groups Q and $Q(\alpha)$ actually have the desired order q^{d-1}.* This is somewhat similar to what occurred in **[KS]**, 5.3.1. Until we have verified this in Section 2.8 we merely *assume* that these groups are correct. We have seen that this assumption is correct with high probability.

When $d = 4$ we proceed more simply, as well as deterministically, by defining $Q = \langle X_1, X_1^\tau \rangle$, $Q(\alpha) = \langle X_2, X_2^{\tau^{j(\gamma)-1}} \rangle$ and $Q(\gamma) = Q(\alpha)^{j(\gamma)}$. Since the methods for constructing X_1, X_2, τ and $j(\gamma)$ differ from **[KS]** in this case, we supply a proof that Q and $Q(\alpha)$ behave as desired.

Let $\Psi : \mathrm{SL}(4, q) = \mathrm{SL}(V) \to G$ denote the target isomorphism and identify V_J with the 3-dimensional support of $\Psi^{-1}(J)$ in V. Then $\Psi^{-1}(T) = T_{\alpha, W}$ where $W \cap V_J = H$. Let δ denote the 1-space of V centralised by $\Psi^{-1}(J)$, so that $W = \langle H, \delta \rangle$ and δ is in the 3-dimensional eigenspace of $\Psi^{-1}(a)$. Hence $\Psi^{-1}(\tau)$ acts irreducibly on W and $\langle X_1, X_1^\tau \rangle$ is the image under Ψ of the group of transvections with axis W. Also, $\Psi^{-1}(\tau)$ fixes γ so that $(\Psi^{-1}(\tau))^{(B')^{-1}}$ fixes α: $Q(\alpha)$ is the image under Ψ of the group of transvections with centre α.

The subgroup H. Let $H = \langle (X_1D)^{\sigma_1^i}, (X_1D)^{(\sigma_2^f)^i} \mid 0 \le i \le d \rangle$ if $d > 4$ and $\langle X_1D, (X_1D)^\tau \rangle$ if $d = 4$. The argument in **[KS]**, 3.3.1, shows that $H = N_G(Q)'$ *provided that Q is the correct group,* of order q^{d-1} (as is already the case when $d = 4$). Thus, $H = N_G(Q)'$ with high probability (and this is verified in Section 2.8).

Transitivity of Q and $Q(\alpha)$. The following variation of **[KS]**, Lemma 3.10, is not deterministic: the restriction imposed by our timing goals seems to force a departure from **[KS]** in this regard. While our use of randomization only takes place inside elementary abelian groups, we cannot be sure of the correctness of the output of this lemma until we are sure of the correctness of Q (cf. Section 2.8).

Note that **[KS]**, 3.3.2, contains simple and fast procedures for deciding whether or not two conjugates $Q(\alpha)^x$ and $Q(\alpha)^y$ (or Q^x and Q^y) are equal, or whether the "point" $Q(\alpha)^x$ is "on" the "hyperplane" Q^y. There is also a simple test for whether two conjugates T_1, T_2 of T are opposite (i.e., generate an $\mathrm{SL}(2, q)$): $[[t_1, t_2,]t_2] \ne 1$ for just one choice of $1 \ne t_i \in T_i$.

Lemma 2.1. *There are $O(\xi + \chi e + \mu(d + \log^2 q))$-time Las Vegas algorithms which, with probability $\ge \frac{1}{2^6}$, solve the following problems.*

(i) *Given conjugates $Q(\alpha)^x$ and $Q(\alpha)^y$ not on Q, find the unique $u \in Q$ such that $Q(\alpha)^{xu} = Q(\alpha)^y$.*

(ii) *Given conjugates Q^x and Q^y not on $Q(\alpha)$, find the unique $u \in Q(\alpha)$ such that $Q^{xu} = Q^y$.*

Proof. We only prove (i). We may assume that $Q(\alpha)^x \ne Q(\alpha)^y$. Let $T_0 = T$ and let T_1 be one of the conjugates of T generating $Q(\alpha)^x$ and opposite to T. Choose (generators

for) a random transvection group $T_2 < Q(\alpha)^y$, use Section 3 to test whether $K = \langle T_0, T_1, T_2 \rangle \cong \mathrm{SL}(3, q)$, and if so to find an isomorphism $\Psi_K \colon \mathrm{SL}(3, q) \to K$. Use Ψ_K to find the subgroup Q_{Ki} of K that contains T_i, consists of q^2 transvections (for $i = 0, 1, 2$), and satisfies $[Q, Q_{K0}] = [Q(\alpha)^x, Q_{K1}] = [Q(\alpha)^y, Q_{K2}] = 1$. Finally, use Lemma 3.1 to find the element $u \in Q_{K0}$ such that $Q_{K1}^u = Q_{K2}$. Output u.

We claim that u behaves as required with the stated probability. First note that $K \cong \mathrm{SL}(3, q)$ with probability $\geq \frac{1}{2^5}$ (as in **[KS]**, Lemma 3.7), while Section 3 verifies this with probability $\geq \frac{1}{2}$, so we obtain K behaving as stated, together with the needed isomorphism, with probability $\geq \frac{1}{2^6}$. Since K acts on the target vector space V by preserving a decomposition as a direct sum of a 3- and a $d - 3$-space, each Q_{Ki} consists of transvections of V, $Q_{K1} \leq Q(\alpha)^x$, $Q_{K2} \leq Q(\alpha)^y$ and $u \in Q_{K0} \leq Q$. The timing is dominated by the use of the above commutator tests and the use of Theorem 1.1(1) with $d = 3$. □

We note that the corresponding time in **[KS]** was $O(\mu(qe+qd))$ for a *deterministic* algorithm. However, we will always need to repeat the lemma at least $\lceil 2^6 \log d \rceil$ times in order to ensure that it produces an output with probability $> 1 - \frac{1}{d^2}$; we will see this very soon below, as well as later in Section 2.8. We note also that the probability $\frac{1}{2^6}$ can be increased considerably here, since we are generating with full transvection groups, and we are assuming that $q \geq 17$.

The subgroup L. Next we (probably) find a subgroup L of G such that $N_G(Q)' = Q \rtimes L$, essentially as in **[KS]**, Corollary 3.12. For each choice of $i = 1, 2$ and $0 \leq j \leq d - 2$, use up to $\lceil 2^6 \log(16d^3) \rceil$ calls to Lemma 2.1 to find $u_{i,j} \in Q$ such that $Q(\gamma)^{\sigma_i^j u_{i,j}} = Q(\gamma)$. Let $L = \langle D^{\sigma_i^j u_{i,j}} \mid i = 1, 2; \ 0 \leq j \leq d - 2 \rangle$.

That L behaves as stated is clear (cf. **[KS]**, Corollary 3.12). For each i, j, with probability $\geq 1 - \frac{1}{16d^3}$ at least one of our calls succeeds, and hence all $u_{i,j}$ are found with probability $\geq 1 - \frac{1}{8d^2}$. The procedure takes $O(d \log d(\xi + \chi e + \mu d \log^2 q))$ time in view of the $\Theta(d \log d)$ calls to Lemma 2.1. □

2.6. Recursion. Recursively we find an isomorphism $\Psi_L \colon \mathrm{SL}(d - 1, q) = \mathrm{SL}(V_{d-1}) \to L$; however, as in Section 2.3 and 2.5 we do not include the verification part of the algorithm in this call, postponing verification until Section 2.8. If this recursive call fails then the algorithm terminates. Failure can occur for two very different reasons: a group Q, $Q(\alpha)$ or L constructed in Section 2.5 was not the desired group; or these groups were correct but the recursive call failed to give an output. The probability of any of these events is small, and hence we can ignore them: we have an isomorphism Ψ_L.

We may need to modify the isomorphism Ψ_L in order to ensure that the stabiliser of a 1-space in the natural $\mathrm{SL}(d - 1, q)$-module is sent to the stabiliser of a 1-space of Q; see **[KS]**, 3.3.3.

2.7. Linear algebra. We now need analogues of **[KS]**, Lemma 3.13 and Section 3.4.3, which deal with workable bases for Q and for $Q(\alpha)$. Proceeding as in **[KS]**, Lemma 3.13, we obtain a direct sum decomposition $Q = A_1 \oplus \ldots \oplus A_{d-1}$; we obtain (using Ψ_L) an element $c \in L$ such that $A_i = A_1^{c^{i-1}}$ for $1 \leq i \leq d-1$; and we have a deterministic $O(\mu d)$ algorithm to express any given $u \in Q$ in the form $u = \prod_1^{d-1} a_i$ with $a_i \in A_i$. While **[KS]**, 3.4.2, goes through without additional change, we need replacement routines for those found in **[KS]**, 3.4.3, since we wish to avoid storing an entire transvection group. This is easily achieved, with the aid of our hypothesised oracle, by noting that A_1 lies inside an $\mathrm{SL}(2, q)$ subgroup L_1 of of J, where we already have a field to work with.

As in **[KS]**, 3.4.3(1), let $\mathcal{B} = \{b_i\}_{i=1}^{d-1}$ be an $\mathbb{F}$-basis for Q, where $b_i = b_1^{c^{i-1}}$ for $1 \leq i \leq d-1$. The $\mathrm{SL}(2, q)$ oracle gives an $\mathbb{F}_p$-basis for A_1, namely the image of $\left\{ \begin{pmatrix} 1 & 0 \\ \rho^k & 1 \end{pmatrix} \mid 0 \leq k < e \right\}$, where we may assume that b_1 corresponds to $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. We now obtain the $\mathbb{F}_p$-basis $\mathcal{B}_p$ of **[KS]**, 3.4.3(1), as the union of conjugates of the $\mathbb{F}_p$-basis of A_1 by the various powers of c.

Next, if we are given $u \in Q$ as in **[KS]**, 3.4.3(4), then use **[KS]**, Lemma 3.13 to write $u = \prod_{i=1}^{d-1} a_i$, with $a_i \in A_i$. For each $1 \leq i \leq d-1$, compute $a_i^{c^{1-i}} \in A_1$ and find the $\lambda_i \in \mathbb{F}$ such that the oracle for L_1 assigns $\begin{pmatrix} 1 & 0 \\ \lambda_i & 1 \end{pmatrix}$ to $a_i^{c^{1-i}}$. The coefficient of each b_i can therefore be found in $O(\chi + \mu)$ time, so that the $\mathbb{F}$-vector of u relative to $\mathcal{B}$ is found in $O(d(\chi + \mu))$ time. It is now an easy matter to obtain an SLP of length $O(d \log q)$ from $\mathcal{B}_p$ to u. The time for **[KS]**, 3.4.3(4), is now $O(\chi d + \mu d \log q)$ compared with $O(\mu q d)$ in **[KS]**.

Note that the same procedures apply even when $d = 3$ once we have obtained Q and $Q(\alpha)$ together with suitable direct sum decompositions.

2.8. Conclusion. As in **[KS]**, 3.4.4, we extend Ψ to QL, but here we use Section 2.7 for linear algebra in Q. We label conjugates of $Q(\alpha)$ as in **[KS]**, 3.5.1, replacing **[KS]**, Lemma 3.13, with our Lemma 2.1. Hence, we label a single given conjugate $Q(\alpha)^x$, with probability at least $\frac{1}{2^6}$ in time $O(\xi + \chi e + \mu(d + \log^2 q))$.

We next define $O(d^2 \log q)$-element generating sets $\mathcal{S}^*$ of G and X of $\mathrm{SL}(d, q)$ and a bijection $X \to \mathcal{S}^*$ that extends to an epimorphism $\Psi \colon \mathrm{SL}(d, q) \to G$, as in **[KS]**, 3.5.2. Computing $\Psi^{-1}(\mathcal{S}^*)$ goes through without change (note that in order to compute $\Psi^{-1}(g)$ for each of the $O(e)$ elements $g \in \mathcal{S}^* \cap J$ we must find the labels of $d + 1$ conjugates of $Q(\alpha)$). We require that *all* of the images are correct with probability $\geq 1 - \frac{1}{16d^2}$, so we repeat Lemma 2.1 $\lceil 2^6 \log(16ed^3) \rceil$ times in each of our $O(ed)$ applications of 3.5.1. Hence, the total time for **[KS]**, 3.5.2, is now $O(d \log(d \log q)\{\xi + \mu d^2 \log q + \chi d \log q\})$.

As part of our recursion we call the subroutines of this and previous Sections $O(d)$ times. It follows that, if we omit the verification of a presentation for G (see below, as well as Section 2.6), then the algorithm runs in time $O(d^2 \log q \log(d \log q)\{\xi + \mu d^2 \log^2 q + \chi d \log q\})$.

We also need procedures for finding SLPs from $\mathcal{S}^*$ to any given element of G, or from X to any given element of SL(d, q). The latter is already in $O(d\chi + \mu d^2 \log q)$ time in **[KS]**, Proposition 3.17, modulo the use of discrete logs, as indicated above in **(v)**; while the following is essentially **[KS]**, Proposition 3.18:

Proof of Theorem 1.1(2b). Use the proof of Proposition 3.18 in **[KS]**, substituting our procedures wherever the timing there involves q. As above, the only care required concerns probability. Namely, the argument in **[KS]** uses Lemma 2.1 to find up to 3 elements of Q or $Q(\alpha)$ and expresses each using an SLP from $\mathcal{S}^*$; finds the matrix A of an element of H as a linear transformation of the $\mathbb{F}$-space Q by using $d - 1$ calls to Section 2.7; forces A to have determinant 1 by using a discrete log call; and applies Theorem 1.1(2a) to A. We repeat Lemma 2.1 $\Theta(\log d)$ times in order to ensure the stated probability, so that the algorithm takes $O(\log d\{\xi + \chi(e+d^2) + \mu d \log q(\log q + d)\} + d^5 \log q)$ time. $\qquad \square$

Ψ is an epimorphism. Finally, we need to confirm this by checking the relations for a suitable short presentation inside G for a quotient of SL(V), which takes time polynomial in all of the preceding procedures (cf. **[KS]**, 7.2.2); and to check that $G = \langle \mathcal{S} \rangle$ is $\langle \mathcal{S}^* \rangle$, by checking that each member of the original generating set $\mathcal{S}$ is in $\langle \mathcal{S}^* \rangle$. We follow the timing calculation in **[KS]**, 7.2.2, keeping in mind that the probability of success of each Las Vegas routine must be large enough so that the probability that any routine fails is small: the total time is $O(\{ed^2 \log(ed)\{\xi + \chi(e + d) + \mu d^2 \log^2 q\} + |\mathcal{S}| \log(d|\mathcal{S}|)\{\xi + \chi(d^2+e) + \mu d \log q(d + \log q) + d^5 \log q\}$, where e is $O(\log q)$. (Note that we repeat the algorithm in Theorem 1.1(2b) $\Theta(\log(d|\mathcal{S}|))$ times for each element of $\mathcal{S}$ in order to make the probabilities behave correctly.) $\quad \square$

3. PSL($3, q$)

Throughout this section we assume that $G \cong$ PSL($3, q$) or SL($3, q$) for known $q \geq 17$. The occurrences of factors q in the timing in **[KS]**, 3.6.3, are as follows:

(**i**) testing isomorphism with SL($2, q$) and finding such an isomorphism;

(**ii**) performing a test using all conjugates of the form A^b, $b \in B$, for a transvection group B and an element or subgroup A;

(**iii**) listing a transvection group, primarily in order to perform linear algebra within that group; or

(**iv**) a discrete log computation in GF(q)*.

Occurrences of these in **[KS]**, 3.6.3, resemble those indicated in Section 2.

Finding τ, a, A and $L \cong$ SL($2, q$). As in **[KS]**, 3.6.3, find an element τ whose order is divisible by a ppd$^\#(p; 2e)$- and a ppd$^\#(p; e)$-number, as well as by a ppd$^\#(p; e/2)$-number if e is even. Let $a = \tau^{2(q+1)}$.

Choose a conjugate b of a, let $A = \langle a, b \rangle$, find $L = A'$ using **[BCFLS]**, and use the hypothesised SL($2, q$)-oracle to test whether $L \cong$ SL($2, q$). If this isomorphism

holds, we also obtain a field $\mathbb{F} = \mathrm{GF}(q)$, an isomorphism $\Psi_L \colon \mathrm{SL}(2, q) \to L$ and a generating set $\mathcal{S}_L^*$ for L.

As in [KS], Lemma 3.8, with probability $\geq \frac{1}{2}$ we have $A \cong \mathrm{SL}(2, q) \circ \langle a \rangle$; finding L uses a Monte Carlo algorithm that is made Las Vegas by the $\mathrm{SL}(2, q)$-oracle call. Finding a, A and L takes $O(\xi + \mu \log^2 q + \chi)$ time.

Elements z, $j(\gamma) \in A$. We know that A is a central product $A = L \circ \langle z \rangle$, and we have found L. We next construct $\langle z \rangle$.

Use our $\mathrm{SL}(2, q)$-oracle to find the two transvection groups T_1, T_2 in L normalised by a, as well as $h \in L$ of order $q - 1$ normalising T_1 and T_2. Note that a has order dividing $(q - 1)/(2, q - 1)$, while h induces an element of order $(q - 1)/(2, q - 1)$ on T_1. Use the discrete log oracle inside L to find an integer n such that $0 \leq n < q - 1$ and h^n and a induce the same scalar on T_1, and set $z = ah^{-n}$. Then $A = L \circ \langle z \rangle$, where z is found in $O(\mu e \log q + \chi)$ time.

Also find a transvection $j(\gamma) \in L$ such that $T_1^{j(\gamma)} = T_2$.

Finding Q and $Q(\alpha)$. Find $Q(\alpha) = \langle T_1, T_1^{\tau} \rangle$ and $Q = \langle T_1, T_1^{\tau^{j(\gamma)^{-1}}} \rangle$ in $O(\mu \log q)$ time. Since τ fixes the eigenspaces of a (on the vector space V we have not yet constructed and hence cannot use for algorithmic purposes), up to duality these are the groups of order q^2 consisting of all transvections having the same centre or axis as T_1, respectively.

The transvection groups $T_1, \ldots, T_6$. We already have transvection groups T_1 and T_2. Rename T_2 as T_4. In $O(\mu e)$ time find $T_2 = [Q(\alpha), z]$, $T_3 = T_2^{j(\gamma)}$, $T_6 = [Q, z]$ and $T_5 = T_6^{j(\gamma)}$.

In order to understand the behavior of these six groups on V, let T_1 and T_4 have centres α and γ, respectively, and let β denote the point of intersection of the axes of T_1 and T_4. Then $T_1 = T_{\alpha, \langle \alpha, \beta \rangle}$ "is" the group of transvections with centre α and axis $\langle \alpha, \beta \rangle$, $T_2 = T_{\alpha, \langle \alpha, \gamma \rangle}$, $T_3 = T_{\gamma, \langle \alpha, \gamma \rangle}$, $T_4 = T_{\gamma, \langle \beta, \gamma \rangle}$, $T_5 = T_{\beta, \langle \beta, \gamma \rangle}$ and $T_6 = T_{\beta, \langle \alpha, \beta \rangle}$. We have $Q = T_1 \oplus T_6$ and $Q(\alpha) = T_1 \oplus T_2$.

Redefine $L = \langle T_2, T_5 \rangle \cong \mathrm{SL}(2, q)$. This normalises Q and $Q(\gamma) = Q(\alpha)^{j(\gamma)}$.

Decompositions of Q and $Q(\alpha)$; linear algebra. We can omit [KS], Corollary 3.12, since we already have L. We may assume that the oracle for L returns Ψ_L mapping $\left\{ \left(\begin{smallmatrix} 1 & 0 \\ \rho^k & 1 \end{smallmatrix} \right) \mid 0 \leq k < e \right\}$ and $\left\{ \left(\begin{smallmatrix} 1 & \rho^k \\ 0 & 1 \end{smallmatrix} \right) \mid 0 \leq k < e \right\}$ into T_2 and T_5 respectively. Let $c = \Psi_L \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix} \right)$, and proceed as in Section 2.7.

Transitivity of Q and $Q(\alpha)$. We require a version of Lemma 2.1 for the case $d = 3$. In the context of this paper, we view the following algorithm as deterministic: the only use of randomness is that which is occurs in any practical algorithm to recognise $\mathrm{SL}(2, q)$.

Lemma 3.1. *There is a deterministic $O(\chi + \mu \log q)$ time algorithm to solve problems* (i) *and* (ii) *of Lemma* 2.1 *when $d = 3$.*

Proof. We redefine $h = \Psi_L(\text{diag}(\rho^{-1}, \rho))$ and view it as acting on the transvection groups T_2 and T_5 as the scalars ρ^{-2} and ρ^2, respectively. It follows (by considering an appropriate isomorphism $\Psi \colon \text{SL}(3, q) \to G$ extending Ψ_L) that we may view h as acting on T_1 and T_6 as the scalars ρ^{-1} and ρ, respectively.

In (i) we are given $Q(\alpha)^x$ and $Q(\alpha)^y$, each generated by two transvection groups (e.g., T_1^x and T_2^x) at least one of which is opposite T_1. Let $A_1 < Q(\alpha)^x$ and $A_2 < Q(\alpha)^y$ be transvection groups opposite T_1. For $i = 1, 2$, use the SL$(2, q)$-oracle for $\langle T_1, A_i \rangle$ to find an element h_i of order $q - 1$ normalising T_1 and A_i. Each h_i induces an automorphism of order $(q - 1)/(2, q - 1)$ on T_1; use the discrete log oracle inside $\langle T_1, A_1 \rangle$ to arrange for these to be the same automorphism as that induced by h^{-2} (which we know corresponds to the scalar ρ^2).

If q is even or $q \equiv 3 \pmod 4$, then ρ is the only square root of ρ^2 in $\mathbb{F}^*$ of order $q - 1$, and hence h_1 and h_2 induce automorphisms on U/T_1 corresponding to ρ. If $q \equiv 1 \pmod 4$ then the automorphisms of order $q - 1$ that h_1 and h_2 induce on U/T_1 correspond to $\pm\rho$, and we now ensure that they both induce the same automorphism ρ. Fix $1 \neq t \in T_6$. For $i = 1, 2$, use the $\langle T_1, A_1 \rangle$-oracle to test whether $b_i = t^{h_i}(t^h)^{-1} \in T_1$; if not then replace h_i by $h_i^{(q+1)/2}$ (since in this situation h_i induces $-\rho$ on U/T_1). Each h_i now induces ρ on U/T_1.

We now construct $u \in U = Q(\alpha)Q = T_2 T_1 T_6$ such that $h_1^u \equiv h_2 \pmod{T_1}$. Note that $h_1^{-k} h_2^k$ centralises both T_1 and U/T_1 for any integer k, and hence lies in U. If q is odd set $u = (h_1^{(q-1)/2} h_2^{(q-1)/2})^{(p+1)/2}$; if q is even set $u = (h_1^{-1} h_2)^{h_1^j}$ for the integer j such that $0 \leq j \leq q - 1$ and $\rho^j(1 - \rho) = 1$, found using the discrete log oracle.

We claim that $h_1^u \equiv h_2 \pmod{T_1}$. Since $U\langle h_1 \rangle/T_1 = U\langle h_2 \rangle/T_1$, there is some $\tilde{u} \in U$ such that $h_1^{\tilde{u}} \equiv h_2 \pmod{T_1}$. We will compute using actions on the group $U/T_1 \cong \mathbb{F}^2$. Suppose that the element $\tilde{u}$ we seek acts on $\mathbb{F}^2$ via $v \mapsto v + c$ for some $c \in \mathbb{F}^2$. Since h_1 acts via $v \mapsto \rho v$ we have $h_2 = h_1^{\tilde{u}} \colon v \mapsto \rho v + (1 - \rho)c$. If q is even then $(h_1^{-1} h_2)^{h_1^j} \colon v \mapsto v + c$, so that $u = (h_1^{-1} h_2)^{h_1^j}$ behaves as claimed; the case q odd is similar.

Next, we decompose $u = v'v$ with $v' \in Q(\alpha) = T_2 T_1$ and $v \in Q = T_6 T_1$. Use our SL$(2, q)$-oracle to find an element $h_{36} \in \langle T_3, T_6 \rangle$ of order $q - 1$ normalising T_3 and T_6. This element induces a scalar of order $q - 1$ on T_2; use the discrete log oracle for L to arrange for this scalar to be ρ^{-1}. It follows (again by considering an extension Ψ of Ψ_L) that h_{36} induces ρ^2 on T_6. Use the discrete log oracle to find integers m and n such that $0 \leq m, n \leq q - 1$, $\rho^{-m}(\rho^{-1} - 1) = 1$ and $\rho^n[\rho^{2m}(\rho^2 - 1) - 1] = 1$. We claim that $v = (u^{-1}[u, h_{36}]^{h_{36}^m})^{h^n}$ satisfies the conditions stated above.

We know that there are elements $v' \in T_2 T_1$ and $v \in T_6 T_1$ such that $u = v'v$. Computing once again inside the 2-space $U/T_1 \cong T_2 \oplus T_6$ we find, mod T_1, first that $[u, h_{36}] \equiv (v' + v)^{h_{36}} - (v' + v) \equiv (\rho^{-1} - 1)v' + (\rho^2 - 1)v$, then that $u^{-1}[u, h_{36}]^{h_{36}^m} \equiv v(\rho^{2m}(\rho^2 - 1) - 1)$, and finally that $(u^{-1}[u, h_{36}]^{h_{36}^m})^{h^n} \equiv v$. Thus, our choice of v behaves as claimed.

Finally, observe that $\langle A_1^u, T_1 \rangle = \langle A_2, T_1 \rangle$ (since $\langle h_i, T_1 \rangle$ is in just one conjugate of L for $i = 1, 2$), and hence $\langle A_2, T_1, A_1^v \rangle/O_p(\langle A_2, T_1, A_1^v \rangle) \cong \text{SL}(2, q)$ (since v' fixes

the line joining α and α^{xu} in the target module V for G). Use our $\mathrm{SL}(2, q)$-oracle for this quotient group to find $t \in T_1$ such that $A_1^{vt} \equiv A_2 \pmod{O_p(\langle A_2, T_1, A_1^v \rangle)}$, and output vt. All of this takes $O(\chi + \mu \log q)$ time. (N.B.—We will use a similar, but simpler, method below in Section 4.2.)   □

Remark. In the proof above we used our $\mathrm{SL}(2, q)$-oracle inside the black box group $K/O_p(K)$ where $K = \langle A_2, T_1, A_1^v \rangle$. In order to do this, we need a membership test for $O_p(K)$ (so that in $K/O_p(K)$ we can decide whether or not a given string represents the identity, and hence $K/O_p(K)$ is a black box group). For any black box group K there is an efficient test for whether a given $k \in K$ lies in $O_p(K)$: test whether or not $\langle k^K \rangle$ is a p-group (see **[BCFLS]**).

End of proof. Now we have all of the subgroups needed in order to complete the case $d = 3$ by direct imitation of **[KS]**, 3.6.3. The timing calculations in **[KS]**, 3.6.3, 7.2.2, show that the time for Theorem 1.1(1) is $O(\xi + e\chi + \mu \log^2 q)$ without the verification of a presentation, and $O(\xi + e\chi + \mu \log^2 q + |\mathcal{S}|(\chi + \mu \log q))$ with such a verification. The times for Theorem 1.1(2a) and (2b) are both $O(\chi + \mu \log q)$.

Note that the use of randomness in Theorem 1.1(2b) when $d > 3$ is forced upon us by the Las Vegas Lemma 2.1. When $d = 3$, however, the analogous Lemma 3.1 above is deterministic, upgrading Theorem 1.1(2b) to a deterministic algorithm in this case. This also accounts for the apparent discrepancy between the timings for $d = 3$ and $d = 4$ in Theorem 1.1(1).

However, as already noted at the end of Section 1, the existing algorithms for recognition of $\mathrm{SL}(2, q)$ involve the use of randomness, so, in practice, probabilities will need to be treated more carefully for $\mathrm{SL}(3, q)$ and $\mathrm{PSL}(3, q)$.

4. Other classical groups

We will now briefly outline an algorithm for a version of Theorem 1.1 for symplectic groups $\mathrm{PSp}(2m, q)$ and $\mathrm{Sp}(2m, q)$ when $2m \geq 4$ and $q > 9$; and then comment on the remaining classical groups.

4.1. PSp$(2m, q)$, $2m \geq 6$. The group $G \cong \mathrm{PSp}(2m, q)$ or $\mathrm{Sp}(2m, q)$, $2m \geq 6$, is actually somewhat simpler to handle than the case considered in Section 2. In **[KS]** the only places a factor of q arises in timings are of five types:

(**i**) testing for isomorphism to $\mathrm{Sp}(4, q)$, $\Omega^+(4, q)$ or $\Omega^+(6, q)$ and finding such an isomorphism;

(**ii**) testing whether a given subgroup consisting of transvections is a full transvection group of order q;

(**iii**) performing a test using all conjugates a^b, $b \in B$, for a short group B and an element a;

(iv) listing a short root group of order q in order to perform linear algebra within that group or to perform a discrete log call; or

(v) listing a transvection group of order q in order to perform linear algebra within that group.

Type (i). These occur in **[KS]**, 5.2.1 and 5.2.2. For $\mathrm{Sp}(4,q)$ see Section 4.2; for $\Omega^+(4,q)$ use **[KS]**, 3.6.2, together with our $\mathrm{SL}(2,q)$-oracle; and for $\Omega^+(6,q)$ use Section 2 since $\mathrm{P}\Omega^+(6,q) \cong \mathrm{PSL}(4,q)$.

Type (ii). This occurs in **[KS]**, 5.3.1. If T denotes the given group, let $g \in G$ and use our $\mathrm{SL}(2,q)$-oracle to test whether $\langle T, T^g \rangle \cong \mathrm{SL}(2,q)$; if so use the resulting isomorphism to test whether $|T| = q$. This isomorphism occurs with high probability, but $\Theta(\log m)$ choices of g are needed to make the probability of success $> 1 - \frac{1}{8m^2}$.

Type (iii). This occurs in **[KS]**, 5.3.2. The context is as follows: at that point there is a group Q that (probably) is $O_p(N_G(T))$ for a transvection group T. The goal is a procedure which, when given conjugates T_1, T_2 of T, neither of which commutes with T, produces the unique element $u \in Q$ conjugating T_1 to T_2. We assume that T is not in $\langle T_1, T_2 \rangle$ (as otherwise we can use our oracle). Then we obtain u as follows: for a random $v \in Q$ use Section 4.2 to test whether $K = \langle T, T_1, T_2, T_1^v \rangle \cong \mathrm{Sp}(4,q)$; and if so, to find $u \in O_p(N_K(T)) \leq O_p(N_G(T))$ conjugating T_1 to T_2 (see Section 4.2, **Type (ii)**). Once again $\Theta(\log m)$ choices of g are needed to make the probability of success $> 1 - \frac{1}{8m^2}$.

Type (iv). This occurs in **[KS]**, 5.4.2 and 5.4.3, and needs to be handled for just one subgroup A: all others that need to be considered are given as explicit conjugates of one of them. This subgroup A may be assumed to lie in a subgroup $\mathrm{Sp}(4,q)$ that has already been constructed, in which case Section 4.2 can be applied. Our oracle also handles the discrete log call appearing in **[KS]**, 5.4.2.

Type (v). This occurs in **[KS]**, 5.4.4, and is handled by our $\mathrm{SL}(2,q)$-oracle.

Proposition 5.18 in [KS]. This contains an algorithm for finding an SLP to any given element $g \in G$, and uses some of the previous procedures whose timings had factors of q.

End of proof. We can now complete the algorithm by direct imitation of **[KS]**, Sections 5 and 7.2.2.

4.2. $\mathrm{PSp}(4,q)$. When $G \cong \mathrm{Sp}(4,q)$ or $\mathrm{PSp}(4,q)$, the places a factor of q appears in the timing in **[KS]**, 5.6.1, are as follows:

(i) testing isomorphism with $\mathrm{SL}(2,q)$ and finding such an isomorphism;

(ii) performing a test using all conjugates a^b, $b \in B$, for a short group B and a known element a;

(iii) listing a short root group of order q in order to perform linear algebra within that group or to perform a discrete log call; or

(iv) listing a transvection group of order q in order to perform linear algebra within that group.

Type (i). This situation, in which our $SL(2, q)$-oracle can be used, occurs when finding a subgroup $A \circ B$ isomorphic to the central product $SL(2, q) \circ SL(2, q)$, as well as when finding $Q = O_p(N_G(T))$ for a transvection group $T < A$.

Type (ii). The context is as in Section 4.1**(iii)**: we need a procedure which, when given conjugates T_1, T_2 of T, neither of which commutes with T, produces the unique element $v \in Q$ conjugating T_1 to T_2. For this purpose, for $i = 1, 2$ apply our oracle to $\langle T, T_i \rangle$ in order to obtain elements h_i that normalise T and T_i, where $|h_i| = 2$ if q is odd and $|h_i| = q - 1$ if q is even. We will construct $u \in Q$ such that $h_1^u \equiv h_2$ (mod T), and hence $\langle T, T_1 \rangle^u = \langle T, T_2 \rangle$, in which case our $SL(2, q)$-oracle can be used to find $t \in T < Q$ such that $T_1^{ut} = T_2$; then we output $v = ut$. That the element u constructed below behaves as required follows from a straightforward calculation (cf. Section 3).

If q is odd then $u = (h_1 h_2)^{(p+1)/2}$ works. If q is even then h_1 and h_2 induce generators of the same cyclic automorphism group of order $q - 1$ on T; use our discrete log oracle to make them induce the same automorphism ρ, and also to find j such that $\rho^j (1 - \rho) = 1$. Then $u = (h_1^{-1} h_2)^{h_1^j}$ works.

Type (iii). As noted above we have already constructed a subgroup $B < C_G(T)$ and an isomorphism $\Psi_B : SL(2, q) \to B$. Use Ψ_B to find an element $h \in B$ of order $q - 1$, as well as the two transvection groups T_1, T_2 of B normalised by h. Then h acts on the abelian groups $Q_i = [Q, T_i]T$ of order q^2; these contain the short root groups we need to handle without listing. Note that $Q/T = (Q_1/T) \times (Q_2/T)$.

We may assume that $\Psi_B \begin{pmatrix} \rho & 0 \\ 0 & \rho^{-1} \end{pmatrix} = h$. In view of the standard matrix representation of $Sp(4, q)$ we may assume that h induces ρ on Q_1/T. That is, we *define* the action of $\langle \rho \rangle$ on Q_1/T by: $\rho^j (u_1 T) = u_1^{h^j} T$ for all $u_1 \in Q_1$ and $0 \le j < q - 1$; and similarly we define $\rho^j (u_2 T) = u_2^{h^{-j}} T$ for $u_2 \in Q_2$. Thus, we can now perform discrete log calls within Q_i/T.

We use two different methods to find the Q_i/T−coordinates of a given uT in Q/T depending upon the parity of q. When q is odd, this is achieved by constructing an alternating form on the 2-space Q/T using commutators in Q. When q is even, Q is abelian, so we must use a different method. In this case, however, short root groups are $\mathrm{Aut}(G)$−conjugate to transvection groups and we can use the $SL(2, q)$−oracle to perform linear algebra inside short root subgroups of Q_i.

Type (iv). As in Section 4.1**(v)** this is handled by our $SL(2, q)$-oracle.

End of proof. We can now complete the algorithm by direct imitation of **[KS]**, 5.6.1, 7.2.2.

Alternative ending. Once we have obtained the exact structure of QB (recall that $B \cong SL(2, q)$ centralises T), it is not difficult to proceed directly to a short presentation for G, as in **[KM]**.

4.3. Orthogonal and unitary groups. The algorithms in **[KS]** for orthogonal groups involve factors of q in their timings in more ways than in the other classical groups. Nevertheless, there is reasonable hope that these factors can be replaced by calls to our SL$(2, q)$-oracle.

On the other hand, the timing in **[KS]** for PSU(d, q) involves $q^{3/2}$, starting with the 3-dimensional groups. However, it seems likely that it will be possible to prove an analogue of Theorem 1.1 in this case, assuming the availability of a PSU$(3, q)$-oracle. Moreover, we hope that hypothesising such an oracle will be justified for PSU$(3, q)$, given as a group of 3×3 matrices, by a variation on ideas in **[CLG]**.

References

[BCFLS]   L. Babai, G. Cooperman, L. Finkelstein, E. M. Luks and Á. Seress, Fast Monte Carlo algorithms for permutation groups, J. Comput. System Sci. 50 (1995) 296–308.

[Br]   S. Bratus, Recognition of finite black box groups, Ph.D. Thesis, Northeastern Univ. 1999.

[CLG]   M. Conder and C. R. Leedham-Green, Fast recognition of classical groups over large fields, in: Groups and Computation III (W. M. Kantor and Á. Seress, eds.), Ohio State Univ. Math. Res. Inst. Publ. 8, Walter de Gruyter, Berlin–New York, 2001, 113–121.

[KL]   W. M. Kantor and R. A. Liebler, The rank 3 permutation representations of the finite classical groups, Trans. Amer. Math. Soc. 271 (1982) 1–71.

[KM]   W. M. Kantor and K. Magaard, Black box exceptional groups of Lie type, in preparation.

[KS]   W. M. Kantor and Á. Seress, Black box classical groups, to appear in Mem. Amer. Math. Soc.

Department of Mathematics
University of Oregon
Eugene, OR 97403, U.S.A.
brooks@geometry.uoregon.edu
kantor@math.uoregon.edu